

Distributed Functional Scalar Quantization Simplified

John Z. Sun*, *Student Member, IEEE*, Vinith Misra, *Student Member, IEEE*,
and Vivek K Goyal, *Senior Member, IEEE*

Abstract

Distributed functional scalar quantization (DFSQ) theory provides optimality conditions and predicts performance of data acquisition systems in which a computation on acquired data is desired. We address two limitations of previous works: prohibitively expensive decoder design and a restriction to sources with bounded distributions. We rigorously show that a much simpler decoder has equivalent asymptotic performance as the conditional expectation estimator previously explored, thus reducing decoder design complexity. The simpler decoder has the feature of decoupled communication and computation blocks. Moreover, we extend the DFSQ framework with the simpler decoder to acquire sources with infinite-support distributions such as Gaussian or exponential distributions. Finally, through simulation results we demonstrate that performance at moderate coding rates is well predicted by the asymptotic analysis, and we give new insight on the rate of convergence.

I. INTRODUCTION

FUNCTIONAL source coding techniques are of great importance in modern distributed systems such as sensor networks and cloud computing architectures because the fidelity of acquired data can greatly impact the accuracy of computations made with that data. In this work, we provide theoretical and empirical results for quantization in distributed systems described by the topology in Fig. 1. Here, N memoryless sources produce scalar realizations $X_1^N = (X_1, \dots, X_N)$ from a joint distribution $f_{X_1^N}$ at each discrete time instant. These measurements are compressed via separate encoders and then sent to a central decoder that approximates a computation on the original data; the computation may be the identity function, meaning that the acquired samples themselves are to be reproduced.

There has been substantial effort to study distributed coding using information-theoretic concepts, taking advantage of large block lengths and powerful decoders to approach fundamental limits of compression. However, techniques inspired by this theory are infeasible for most applications. In particular, strong dependencies between source variables imply low information content per variable, but exploiting this is difficult under rigid latency requirements.

Rather than have long blocks, the complementary asymptotic of *high-resolution quantization theory* [1] is more useful for these scenarios; most of this theory is focused on the scalar case, where the block length is one. The principal previous work in applying high-resolution quantization theory to the acquisition and computation network of Fig. 1 is the *distributed functional scalar quantization* (DFSQ) framework [2]. The key message from DFSQ is that the design of optimal encoders for systems that perform nonlinear computations can be drastically different from what traditional quantization theory suggests. In recent years, ideas from DFSQ have been applied to compressed sensing [3], compression for media [4], and channel state feedback in wireless networks [5].

Like the information-theoretic approaches, the existing DFSQ theory relies in principle on a complicated decoder. (This reviewed in Section II-C.) The primary contribution of this paper is to study a DFSQ framework that employs a simpler decoder.

This material is based upon work supported by the National Science Foundation under Grants 0643836, 0729069, and 1115159.

J. Z. Sun and V. K. Goyal are with the Department of Electrical Engineering and Computer Science and the Research Laboratory of Electronics, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: {johnsun,vgoyal}@mit.edu).

V. Misra is with the Department of Electrical Engineering and the Information Systems Laboratory, Stanford University, Stanford, CA 94305 USA (e-mail: vinith@stanford.edu).




Fig. 1. A distributed computation network, where each of N spatially-separated sources generate a scalar X_n . The scalars are encoded and communicated over rate-limited links to a central decoder without interaction between encoders. The decoder computes an estimate of the function $g(X_1^N) = g(X_1, X_2, \dots, X_N)$ from the received data using $\hat{g}(\hat{X}_1^N)$. Each encoder is allowed transmission rate R_n .




Fig. 2. A block diagram for companding as a constructive method for non-uniform scalar quantization. The notation $Q_{U,K}$ is used to describe the canonical uniform quantizer with K partitions in the granular region $[0, 1]$. In this paper, only the partition boundaries are scaled using compressor function c ; the codewords are defined through midpoint reconstruction (2) and can be computed at the decoder.

Remarkably, the same asymptotic performance is obtained with the simpler decoder, so the optimization of quantizer point density is unchanged. Furthermore, the simplified framework allows a greater decoupling or modularity between communication (source encoding/decoding) and computation aspects of the network.

The analysis presented here uses different assumptions on the source distributions and function than [2]—neither is uniformly more or less restrictive. Unlike in [2], we are able to allow the source variables to have infinite support. In fact, the functional setting allows us to present high-resolution quantization results for certain heavy-tailed source distributions for the first time.

We begin in Sec. II by reviewing relevant previous work and summarizing the contributions of this paper. In Sec. III and IV, we give distortion and design results for a distributed network. Finally, we provide examples for the theory in Sec. V and conclude in Sec. VI.

II. PRELIMINARIES

A. Previous Work

The distributed network shown in Fig. 1 is of great interest to the information theory and communications communities, and there exists a variety of results corresponding to different scenarios of interest. We present a short overview of some major works; a comprehensive review appears in [2].

In the large block length asymptotic, there are many influential and conclusive results. For the case of discrete-valued sources and $g(X_1^N) = X_1^N$, the lossless distributed source coding problem is solved by Slepian and Wolf [6]. In the lossy case, the problem is generally open except in specific situations [7], [8]. The case where $g(X_1^N) = X_1$ and the rate is unconstrained except for R_1 is the well-known source coding with side information problem [9]. For more general computations, the lossless [10]–[12] and lossy [13], [14] cases have both been explored.

There are also results for when the block length is constrained to be very small. We will delay discussion of DFSQ for later and instead focus on related works. The use of high-resolution for computation has been considered in detection and estimation problems [15]–[17]. In the scalar setting, the scenario where the computation is unknown but is drawn from a set of possibilities has been studied [18]. Finally, there are strong connections between DFSQ and multidimensional companding, a technique used in perceptual coding [19].

B. High-resolution Scalar Quantizer Design

A scalar quantizer Q_K is a mapping from the real line to a set of K points $\mathcal{C} = \{c_k\}_{k=1}^K \subset \mathbb{R}$ called the codebook, where $Q_K(x) = c_k$ if $x \in P_k$ and the cells $\{P_k\}_{k=1}^K$ form a partition of \mathbb{R} . The quantizer is called *regular* if the partition cells are intervals containing the corresponding codewords. We then assume the codebook entries are indexed from smallest to largest and that $P_k = (p_{k-1}, p_k]$ for each k ; this is essentially without loss of generality because the dispositions of the endpoints of the cells are immaterial to performance when the quantizer input is continuous. Regularity implies $p_0 < c_1 \leq p_1 < c_2 \leq \dots < c_K \leq p_K$, with $p_0 = -\infty$ and $p_K = \infty$. Define the *granular* region as (c_1, c_K) and its complement $(-\infty, c_1] \cup [c_K, \infty)$ as the *overload* region.

Uniform (linear) quantization, where partition cells in the granular region have equal length, is most commonly used in practice, but other quantizer designs are possible. Fig. 2 presents the compander model as a method for generating nonuniform quantizers from a uniform one. In this model, the scalar source is transformed using a nondecreasing and smooth *compressor* function $c : \mathbb{R} \rightarrow [0, 1]$, then quantized using a uniform quantizer comprising K levels on the granular region $[0, 1]$, and finally passed through the *expander* function c^{-1} . Compressor functions are defined such that $\lim_{x \rightarrow -\infty} c(x) = 0$ and $\lim_{x \rightarrow \infty} c(x) = 1$. It is convenient to define a *point density function* as $\lambda(x) = c'(x)$. Because of the extremal conditions on c , there is a one-to-one correspondence between λ and c , and hence a quantizer of the form shown in Fig. 2 can be uniquely specified using a point density function and codebook size. We denote such a quantizer as $Q_{K,\lambda}$. By virtue of this definition, the integral of the point density function over any quantizer interval is $1/K$:

$$\int_{p_k}^{p_{k+1}} \lambda(x) dx = \frac{1}{K}, \quad k = 1, 2, \dots, K. \quad (1)$$

In practice, scalar quantization is rarely, if ever, performed by an explicit companding operation. A slight modification that avoids repeated computation of c^{-1} is to apply the compressor c , compare to threshold values (multiples of $1/K$) to determine

the partition cell P_k , and then obtain c_k from a pre-computed table. We assume that the non-extremal reconstruction values are set to the midpoints of the cells, i.e.

$$c_k = \frac{p_{k-1} + p_k}{2}, \quad k = 2, 3, \dots, K-1. \quad (2)$$

This is suboptimal relative to centroid reconstruction, but has the simplicity of depending only on λ and K —not on the source density. The extremal reconstruction values are fixed to be $c_1 = p_1$ and $c_K = p_{K-1}$. This again is suboptimal but does not depend on the source distribution. We will show later that this construction does not affect asymptotic quantizer performance.

The utility of the compander model is that we can precisely analyze the distortion behavior as K becomes large and use this to optimize λ . Assuming the source is well-modeled as being drawn iid from a probabilistic distribution, we define the mean-squared error (MSE) distortion as

$$D_{\text{mse}}(K, \lambda) = \mathbb{E}[|X - Q_{K,\lambda}(X)|^2], \quad (3)$$

where the expectation is with respect to the source density f_X . Under the additional assumption that f_X is continuous (or simply measurable) with tails that decay sufficiently fast,

$$D_{\text{mse}}(K, \lambda) \simeq \frac{1}{12K^2} \mathbb{E}[\lambda^{-2}(X)], \quad (4)$$

where \simeq indicates that the ratio of the two expressions approaches 1 as K increases [20], [21]. Hence, the MSE performance of a scalar quantizer can be approximated by a simple relationship between the source distribution, point density and codebook size, and this relation becomes more precise with increasing K . Moreover, quantizers designed according to this approximation are *asymptotically optimal*, meaning that the quantizer optimized over λ has distortion that approaches the performance of the best Q_K found by any means [22]–[24], meaning

$$\inf_{Q_K} \mathbb{E}[|X - Q_K(X)|^2] \simeq \frac{1}{12K^2} \mathbb{E}[\lambda^{-2}(X)]. \quad (5)$$

Experimentally, the approximation is accurate even for moderate K [1], [25]. Since distortion depends only on λ in the asymptote, calculus techniques can be used to optimize companders.

When the quantized values are to be communicated or stored, it is natural to map codewords to a string of bits and consider the trade-off between performance and communication rate R , defined to be the expected number of bits per sample. In the simplest case, the codewords are indexed and the communication rate is $R = \log_2(K)$; this is called *fixed-rate* or *codebook-constrained* quantization. Hölder's inequality can be used to show that the optimal point density for fixed-rate is

$$\lambda_{\text{mse,fr}}^*(x) \propto f_X^{1/3}(x), \quad (6)$$

and the resulting distortion is

$$D_{\text{mse,fr}}^*(R) \simeq \frac{1}{12} \|f_X\|_{1/3} 2^{-2R}, \quad (7)$$

with the notation $\|f\|_p = (\int_{-\infty}^{\infty} f^p(x) dx)^{1/p}$ [26].

In general, the codeword indices can be coded to produce bit strings of different lengths based on probabilities of occurrence; this is referred to as *variable-rate* quantization. If the decoding latency is allowed to be large, one can employ block entropy coding and the communication rate approaches $H(Q_{K,\lambda}(X))$. This particular scenario, called *entropy-constrained* quantization, can be analyzed using Jensen's inequality to show the optimal point density $\lambda_{\text{mse,ec}}^*$ is constant on the support of the input distribution [26]. The optimal quantizer is uniform and the resulting distortion is

$$D_{\text{mse,ec}}^*(R) \simeq \frac{1}{12} 2^{-2(R-h(X))}. \quad (8)$$

Note that block entropy coding suggests that the sources are transmitted in blocks even though the quantization is scalar. As such, (8) is an asymptotic result and serves as a lower bound on practical entropy coders with finite block lengths that match the latency restrictions of a system.

In general, the optimal entropy-constrained quantizer (at a finite rate) for a distribution with unbounded support can have an infinite number of codewords [27]. The compander model used in this paper cannot generate all such quantizers. A common alternative is to allow the codomain of c to be \mathbb{R} rather than $[0, 1]$, resulting in a point density that cannot be normalized [28], [29]. To avoid parallel developments for normalized and unnormalized point densities, we restrict our attention to quantizers that have a finite number of codewords K at any finite rate R . This may preclude exact optimality, but it does not change the asymptotic behavior as K and R increase without bound. Specifically, the contribution to overall distortion from the overload region is made negligible as K and R increase, so the distinction between having finitely- or infinitely-many codewords becomes unimportant.

C. Functional Scalar Quantizer Design

In a distributed network where the encoders employ scalar quantization and the decoder performs a known computation, optimizing for the computation rather than source fidelity can lead to substantial gains. In [2], distortion performance and quantizer design are discussed for the distributed setting shown in Fig. 1, with g a scalar-valued function. For DFSQ, the cost of interest is functional MSE (fMSE):

$$D_{\text{fmse}}(K_1^N, \lambda_1^N) = \mathbb{E} \left[|g(X_1^N) - \hat{g}(Q_{K_1^N, \lambda_1^N}(X_1^N))|^2 \right], \quad (9)$$

where g is a scalar function of interest, \hat{g} is the optimal fMMSE estimator

$$\hat{g}(x_1^N) = \mathbb{E} \left[g(X_1^N) \mid Q_{K_1^N, \lambda_1^N}(X_1^N) = Q_{K_1^N, \lambda_1^N}(x_1^N) \right], \quad (10)$$

and $Q_{K_1^N, \lambda_1^N}$ is scalar quantization performed on a vector such that

$$Q_{K_1^N, \lambda_1^N}(x_1^N) = (Q_{\lambda_1, K_1}(x_1), \dots, Q_{\lambda_N, K_N}(x_N)).$$

Note the complexity of computing \hat{g} : it requires integrating over an N -dimensional partition cell with knowledge of the joint source density $f_{X_1^N}$. Later in this paper, we avoid this complexity by setting \hat{g} to equal g .

Before understanding how a quantizer affects fMSE, it is convenient to define how a computation locally affects distortion.

Definition 1. The *univariate functional sensitivity profile* of a function g is defined as

$$\gamma(x) = |g'(x)|.$$

The *n th functional sensitivity profile* of a multivariate function g is defined as

$$\gamma_n(x) = (\mathbb{E} [|g_n(X_1^N)|^2 \mid X_n = x])^{1/2}, \quad (11)$$

where $g_n(x)$ is the partial derivative of g with respect to its n th argument evaluated at the point x .

Given the sensitivity profile, the main result of [2] says

$$D_{\text{fmse}}(K_1^N, \lambda_1^N) \simeq \sum_{n=1}^N \frac{1}{12K_n^2} \mathbb{E} \left[\left(\frac{\gamma_n(X_n)}{\lambda_n(X_n)} \right)^2 \right], \quad (12)$$

provided the following conditions are satisfied:

MF1. The function g is Lipschitz continuous and twice differentiable in every argument except possibly on a set of Jordan measure 0.

MF2. The source pdf $f_{X_1^N}$ is continuous, bounded, and supported on $[0, 1]^N$.

MF3. The function g and point densities λ_n allow $\mathbb{E}[(\gamma_n(X_n)/\lambda_n(X_n))^2]$ to be defined and finite for all n .

Following the same recipes to optimize over λ_1^N , the relationship between distortion and communication rate is found. In both cases, the sensitivity acts to shift quantization points to where they can reduce the distortion in the computation. For fixed rate, the minimum high-resolution distortion is achieved by

$$\lambda_{n, \text{fmse}, \text{fr}}^*(x) \propto (\gamma_n(x) f_{X_n}(x))^{1/3}, \quad (13)$$

where f_{X_n} is the marginal distribution of X_n . In the entropy-constrained case, the optimizing point density is

$$\lambda_{n, \text{fmse}, \text{ec}}^*(x) \propto \gamma_n(x). \quad (14)$$

Notice unnormalized point densities are not required here since the source is assumed to have bounded support.

D. Main Contributions of Paper

The central goal of this paper is to develop a more practical method upon the theoretical foundations of [2]. In particular, we provide new insight on how a simplified decoder can be used in lieu of the optimal one in (10). Although the conditional expectations are offline computations, they may be extremely difficult and are computationally infeasible for large N and K . We consider the case when the decoder is restricted to applying the function g explicitly on the quantized measurements. To accommodate this change and provide more intuitive proofs, a slightly different set of conditions is required of g , λ_1^N , and $f_{X_1^N}$.

Additionally, we generalize the theory to infinite-support source variables and vector-valued computations. In brief, we derive new conditions on the tail of the source density and computation that allow the distortion to be stably computed. Interestingly, this extends the class of probability densities under which high-resolution analysis techniques have been successfully applied. The generalization to vector-valued g is a more straightforward extension that is included for completeness. We present several examples to illustrate the framework and the convergence to the asymptotics developed in this work.

III. UNIVARIATE FUNCTIONAL QUANTIZATION

We first discuss the quantization of a scalar random variable X by $Q_{K,\lambda}$, with the result fed into a function g in order to approximate $g(X)$. As mentioned, this is a simpler decoder than analyzed in [2]. We find the dependence of fMSE on λ and then optimize with respect to λ to minimize fMSE.

Assume a companding quantizer with point density λ and granular region $S_K \subset \mathbb{R}$. Consider the following conditions on the computation g and the density f_X of the source:

UF1'. The source pdf f_X is continuous and strictly positive on S_K for any finite K .

UF2'. The function g is continuous on S_K with both $|g'|$ and $|g''|$ defined and bounded by a finite constant C_u .

UF3'. $f_X(x)|g'(x)|^{2-m}/\lambda^{2+m}(x)$ is Riemann integrable over S_K for $m = 0, 1, 2$.

UF4'. f_X , g and λ satisfy the tail condition

$$\lim_{y \rightarrow \infty} \frac{\int_y^\infty |g(x) - g(y)|^2 f_X(x) dx}{\left(\int_y^\infty \lambda(x) dx\right)^2} = 0,$$

and the corresponding condition for $y \rightarrow -\infty$.

The main result of this section is on the fMSE induced by a quantizer $Q_{K,\lambda}$ under these conditions:

Theorem 1. Assume f_X , g , and λ satisfy conditions UF1'–UF4'. Then the fMSE

$$D_{\text{fmse}}(K, \lambda) = \mathbb{E} [|g(X) - g(Q_{K,\lambda}(X))|^2] \quad (15)$$

has the following limit:

$$\lim_{K \rightarrow \infty} K^2 D_{\text{fmse}}(K, \lambda) = \frac{1}{12} \mathbb{E} \left[\left(\frac{\gamma(X)}{\lambda(X)} \right)^2 \right]. \quad (16)$$

Proof: See Appendix A. ■

A. Remarks

1. The fMSE in (16) is the same as in (12). We emphasize that the theorem shows that this fMSE is obtained by simply applying g to the quantized variables rather than using the optimal decoder \hat{g} from (10). Further analysis on this point is given in Sec. III-C.

2. When g is monotonic, the performance (16) is as good as quantizing and communicating $g(X)$ [2, Lemma 5]. Otherwise, the use of a regular quantizer results in a distortion penalty, as illustrated in Example 1.

3. One key contribution of this theorem is the additional tail condition for infinite-support source densities, which effectively limits the distortion contribution in the overload region. This generalizes the class of probability densities for which distortion can be stably bounded using high-resolution approximations [22]–[24]. We will demonstrate this with quantization of a Cauchy-distributed scalar in Example 2.

4. For linear computations, the sensitivity is flat, meaning the optimal quantizer is the same as in the MSE-optimized case. Hence, functional theory will lead to new quantizer designs only when the computation is nonlinear.

5. In the proof of Theorem 1, the first mean-value theorem is used on both f_X and λ , implying these densities are continuous. However, this requirement can be loosened to piecewise-continuous distributions provided the tail conditions still hold and a minor adjustment is made on how partition boundaries are chosen [23]. Rather than elaborating further, we refer the reader to a similar extension in [2, Sec. III-F]. An equivalent argument can also be made for g having a finite number of discontinuities in its first and second derivatives.

6. The theorem assumes that g is continuous and differentiable on the granular region, meaning the sensitivity is positive. However, for explicit regions where $g'(x) = 0$, the use of “don’t care” regions can be used to relax these conditions [2, Sec. VII].

B. Asymptotically Optimal Quantizer Sequences

Since the fMSE of Theorem 1 matches (12), the optimizing quantizers are the same. Using the recipe of Sec. II-B, we can show the optimal point density for fixed-rate quantization is

$$\lambda_{\text{fmse,fr}}^*(x) = \frac{(\gamma^2(x)f_X(x))^{1/3}}{\int (\gamma^2(t)f_X(t))^{1/3} dt}, \quad (17)$$

with distortion

$$D_{\text{fmse,fr}}^*(R) \simeq \frac{1}{12} \|\gamma^2 f_X\|_{1/3} 2^{-2R}. \quad (18)$$

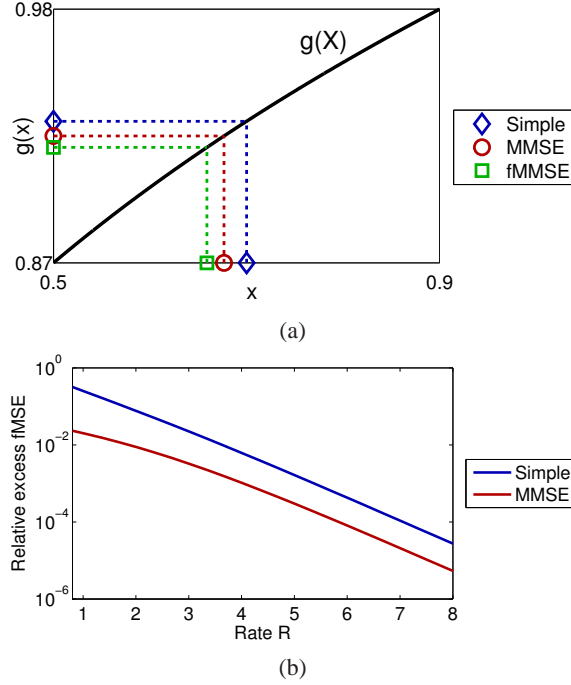


Fig. 3. (a) Codeword placement under simple, MMSE, and fMMSE decoders. The simple decoder performs midpoint reconstruction followed by the application of the computation g . The MMSE decoder applies g to the conditional expectation of X within the cell. Finally, the fMMSE decoder determines (10) for the cell. In this example, the source distribution is exponential and the computation is concave. (b) Performance loss due to the suboptimal codeword placement with respect to rate. We can see that relative excess fMSE decreases linearly with rate and hence the fMSE of the resulting quantizers are asymptotically equivalent.

Meanwhile, optimization in the entropy-constrained case yields

$$\lambda_{\text{fmse,ec}}^*(x) = \frac{\gamma(x)}{\int \gamma(t) dt}, \quad (19)$$

giving distortion

$$D_{\text{fmse,ec}}^*(R) \simeq \frac{1}{12} 2^{2h(X)+2\mathbb{E}[\log \gamma(X)]} 2^{-2R}. \quad (20)$$

C. Negligible Suboptimality of Simple Decoder

Recall the simple decoder analyzed in this work is the computation g applied to midpoint reconstruction as formulated in (2). One can do better by applying g after finding the conditional MMSE estimate of X utilizing knowledge of the source distribution only, or the fMMSE estimator (10) incorporating the function as well. The codeword placements of the three decoders are visualized through an example in Fig. 3(a). The asymptotic match of the performance of the simple decoder to the optimal estimator (10) is a main contribution of this paper.

The simple decoder is suboptimal because it does not consider the source distribution at all, or equivalently assumes the distribution is uniform and the sensitivity is constant over the cell. High-resolution analysis typically approximates the source distribution as uniform over small cells [29], and the proof of Theorem 1 utilizes the fact that the sensitivity is approximately flat over very small regions as well. Hence, the performance gap between the simple decoder and the fMMSE estimator becomes negligible in the high-resolution regime.

To illuminate the rate of convergence, we study the performance gap as a function of quantization cell width, which is dependent on the communication rate (Fig. 3(b)). We see the relative excess fMSE (defined as $(D_{\text{dec}} - D_{\text{opt}})/D_{\text{opt}}$) is exponential in rate, meaning

$$\frac{D_{\text{simple}}}{D_{\text{opt}}} = 1 + c_1 e^{-c_2 R} \quad (21)$$

for some constants c_1 and c_2 . The speed at which the performance gap shrinks contributes greatly to why the high-resolution theory is successful even at low communication rates.

IV. MULTIVARIATE FUNCTIONAL QUANTIZATION

We now describe the main result of the paper for the scenario shown in Fig. 1, where N random scalars (X_1, \dots, X_N) are individually quantized and a scalar computation $g(\hat{X}_1^N)$ is performed. Assume the following conditions on the multivariate joint density, computation and quantizers over a granular region $S_K \subset \mathbb{R}^N$:

MF1'. The joint pdf $f_{X_1^N}$ is continuous and always positive on S_K for any finite K .

MF2'. The multivariate function g is continuous and twice differentiable in every argument over S_K . Every first- and second-order derivative is uniformly bounded by a constant C_m .

MF3'. For any $i, j \in \{1, \dots, n\}$ and $m = 0, 1, 2$, $f_{X_i, X_j}(x_i, x_j) \gamma_i^m(x_i) \lambda_i^{-1}(x_i) \lambda_j^{-1}(x_j)$ is Riemann integrable over S_K .

MF4'. We adopt the notation $x_{\setminus n}$ for x_1^N with the n th element removed; an inverse operator $\tilde{x}(x_n, x_{\setminus n})$ outputs a length- N vector with x_n in the n th element. Then for every index n , the following holds for every $x_{\setminus n}$:

$$\lim_{y \rightarrow \infty} \frac{\int_y^\infty |g(\tilde{x}(x, x_{\setminus n})) - g(\tilde{x}(y, x_{\setminus n}))|^2 f_{X_1^N}(\tilde{x}(x, x_{\setminus n})) dx}{\left(\int_y^\infty \lambda_n(x) dx \right)^2} = 0.$$

An analogous condition holds for the corresponding negative-valued tails.

Recalling $Q_{K_1^N, \lambda_1^N}$ and λ_1^N represent a set of N quantizers and point densities respectively, we present a theorem similar to Theorem 1:

Theorem 2. Assume $f_{X_1^N}$, g , and λ_1^N satisfy conditions MF1'–MF4'. Also assume a fractional allocation α_1^N such that every $\alpha_n > 0$ and $\sum_n \alpha_n = 1$, meaning a set of quantizers $Q_{K_1^N, \lambda_1^N}$ will have $K_n = \alpha_n \kappa$ for some total allocation κ . Then the fMSE

$$D_{\text{fmse}}(K_1^N, \lambda_1^N) = \mathbb{E} \left[|g(X_1^N) - g(Q_{K_1^N, \lambda_1^N}(X_1^N))|^2 \right] \quad (22)$$

of this distributed system has the following limit:

$$\lim_{\kappa \rightarrow \infty} \kappa^2 D_{\text{fmse}}(K_1^N, \lambda_1^N) = \sum_{n=1}^N \frac{1}{12\alpha_n^2} \mathbb{E} \left[\left(\frac{\gamma_n(X_n)}{\lambda_n(X_n)} \right)^2 \right]. \quad (23)$$

Proof: See Appendix B. ■

A. Remarks

1. Like in the univariate case, the simple decoder has performance that is asymptotically equivalent to the more complicated optimal decoder (10).

2. Here, the computation cannot generally be performed before quantization because encoders are distributed. The exception is when the computation is *separable*, meaning it can be decomposed into a linear combination of computations on individual scalars. As a result, the sensitivity is no longer a conditional expectation and quantizer design simplifies to the univariate case, as demonstrated in Example 3.

3. The strict requirements of MF1' and MF2' could potentially be loosened. However, simple modification of individual quantizers like in the univariate case is insufficient since discontinuities may lie on a manifold that is not aligned with the partition lattice of the N -dimensional space. As a result, the error from using a planar approximation through Taylor's Theorem will be $\mathcal{O}(1/\kappa)$, which is no longer negligible. However, based on experimental observations, such as in Example 4, we believe that when these discontinuities exist on a manifold of Jordan measure zero their error may be accounted for. Techniques similar to those in the proofs from [2] could potentially be useful in showing this rigorously.

B. Asymptotically Optimal Quantizer Sequences

As in the univariate case, the optimal quantizers match those in previous DFSQ work since the distortion equations are the same. Using Hölder's inequality, the optimal point density for fixed-rate quantization for each source n (communicated with rate R_n) is

$$\lambda_{n, \text{fmse}, \text{fr}}^*(x) = \frac{(\gamma_n^2(x) f_{X_n}(x))^{1/3}}{\int_{-\infty}^{\infty} (\gamma_n^2(t) f_{X_n}(t))^{1/3} dt}, \quad (24)$$

with fMSE

$$D_{\text{fmse}, \text{fr}}^*(R_1^N) \simeq \frac{1}{12} \sum_{n=1}^N \|\gamma_n^2 f_{X_n}\|_{1/3} 2^{-2R_n}. \quad (25)$$

Similarly, the best point density for the entropy-constrained case is

$$\lambda_{n, \text{fmse}, \text{ec}}^*(x) = \frac{\gamma_n(x)}{\int_{-\infty}^{\infty} \gamma_n(t) dt}, \quad (26)$$

leading to a fMSE of

$$D_{\text{fmse,ec}}^*(R_1^N) \simeq \frac{1}{12} \sum_{n=1}^N 2^{2h(X_n) + 2\mathbb{E}[\log \gamma(X_n)]} 2^{-2R_n}. \quad (27)$$

The rate allocations in (25) and (27) are allowed to vary. Given a total communication rate R , the optimal choice of R_1^N is known [2], [30].

C. Vector-valued Functions

In Theorem 2, we assumed the computation g is scalar-valued. For completeness, we now consider vector-valued functions, where the output of g is a vector in \mathbb{R}^M . Here, the distortion measure is a weighted fMSE:

$$\begin{aligned} D_{\text{fmse}}(K_1^N, \lambda_1^N, \beta_1^M) \\ = \sum_{m=1}^M \beta_m \mathbb{E} \left[|g^{(m)}(X_1^N) - g^{(m)}(Q_{K_1^N, \lambda_1^N}(X_1^N))|^2 \right], \end{aligned} \quad (28)$$

where β_1^M is a set of scalar weights and $g^{(m)}$ is the m th entry of the output of g . Through a natural extension of the proof of Theorem 2, we can find the limit of the weighted fMSE assuming each entry of the vector-valued function satisfies MF1'–MF4'.

Corollary 1. *The weighted fMSE of a source $f_{X_1^N}$, computation g , set of scalar quantizers $Q_{K_1^N, \lambda_1^N}$, and fractional allocation α_1^N has the following limit:*

$$\lim_{\kappa \rightarrow \infty} \kappa^2 D_{\text{fmse}}(K_1^N, \lambda_1^N, \beta_1^M) = \sum_{n=1}^N \frac{1}{12\alpha_n^2} \mathbb{E} \left[\left(\frac{\gamma_n(X_n, \beta_1^M)}{\lambda_n(X_n)} \right)^2 \right], \quad (29)$$

where the sensitivity profile is

$$\gamma_n(x, \beta_1^M) = \left(\sum_{m=1}^M \beta_m \mathbb{E} \left[|g_n^{(m)}(X_1^N)|^2 \mid X_n = x \right] \right)^{1/2}. \quad (30)$$

V. EXAMPLES

In this section, we present examples for both univariate and multivariate functional quantization using asymptotic expressions and empirical results from sequences of real quantizers. The empirical results are encouraging since the convergence to asymptotic limits is fast, usually when the quantizer rate is about 4 bits per source variable. This is because the Taylor remainder term in the distortion calculation decays with an extra κ factor, which is exponential in the rate.

A. Examples for Univariate Functional Quantization

Below we present two examples of functional quantization in the univariate case. The theoretical results follow directly from Sec. III.

Example 1. Assume $X \sim \mathcal{N}(0, 1)$ and $g(x) = x^2$, yielding a sensitivity profile $\gamma(x) = 2|x|$. We consider uniform quantizers, optimal “ordinary” quantizers (quantizers optimized for distortion of the source variable rather than the computation) given in Sec. II-B, and optimal functional quantizers given in Sec. III-B, for a range of rates. The point densities of these quantizers, the source density f_X , and computation g satisfy UF1'–UF4' and hence we utilize Theorem 1 to find asymptotic distortion performance. We also design practical quantizers for a range of R and find the empirical fMSE through Monte Carlo simulations using a random Gaussian source. In the fixed-rate case, theoretical and empirical performance are shown (Fig. 4).

The distortion-minimizing uniform quantizer has a granular region that depends on R , which was explored in [31]. Here, we simply perform a brute-force search to find the best granular region and the corresponding distortion. Surprisingly, this choice of the uniform quantizer performs better over moderate rate regions than the MSE-optimized quantizer. This is because the computation is less meaningful where the source density is most likely and the MSE-optimized quantizer places most of its codewords. Hence, one lesson from DFSQ is that using standard high-resolution theory may yield *worse* performance than a naive approach for some computations. Meanwhile, the functional quantizer optimizes for the computation and gives an additional 3 dB gain over the optimal ordinary quantizer. There is still a loss in using regular quantizers due to the computation being non-monotonic. In fact, if the computation can be performed prior to quantization, we gain an extra bit for encoding the magnitude and thus 6 dB of performance. This illustrates Remark 2 of Sec. III-A.

In the fixed-rate case, the empirical performance approaches the distortion limit described by Theorem 1. The convergence is fast and the asymptotic results predict practical quantizer performance at rates as low as 4 bits/sample.

Example 2. Let a source X be distributed according to the Cauchy distribution centered around 0. This heavy-tail density is special in that the mean and all higher moments are not defined. Hence, it does not satisfy the conditions needed for




Fig. 4. Empirical and theoretical performance for the uniform, ordinary and functional quantizers, as well as the case when the computation is performed first. The distribution is standard normal and $g(x) = x^2$. The distortions are multiplied by 2^{2R} to better indicate the convergence results.




Fig. 5. Empirical and theoretical performance for the uniform, ordinary and functional quantizers for Cauchy f_X and $g(x) = e^{-|x|}$. The distortions are multiplied by 2^{2R} to better indicate the convergence results.

high-resolution theory previously specified in [22]–[24]. However, the functional distortion can be asymptotically determined assuming Condition UF4' is satisfied. The computation $g(x) = \exp(-|x|)$ and the Cauchy density satisfy UF4', and we confirm that experimental results match the theoretical computation of asymptotic distortion in Fig. 5.

B. Examples for Multivariate Functional Quantization

We next provide two examples that follow from the theory of Sec. IV.

Example 3. Let N sources be iid standard normal random variables and the computation be $g(x_1^N) = \sum_n x_n^2$. Since the computation is separable, the sensitivity profile of each source is $\gamma_n(x) = |x|$, and the quantizers are the same as in Example 1. The distortion is also the same, except now scaled by N .

Example 4. Let N sources be iid exponential with parameter $\lambda = 1$ and the computation be $g(x_1^N) = \min(x_1^N)$. In this case, Condition MF2' is not satisfied since there exists $N(N-1)/2$ two-dimensional planes where the derivative is not defined. However, as discussed in the remarks of Theorem 2, we strongly suspect we can disregard the distortion contributions from these surfaces. The overall performance, ignoring the violation of condition MF2', may be analyzed using the sensitivity:

$$\begin{aligned} \gamma_n(x) &= (\mathbb{E}[|g_n(X_1^N)|^2 | X_n = x])^{1/2} \\ &= (\Pr\{\min(X_1^N) = X_1 | X_1 = x\})^{1/2} \\ &= (e^{-\lambda x})^{(N-1)/2}, \end{aligned}$$

where the third line follows from the cdf of exponential random variables.

In Fig. 6, we experimentally verify that the asymptotic predictions are precise. This serves as evidence that MF2' may be loosened.

VI. CONCLUSIONS

In this paper, we have extended distributed functional scalar quantization to a general class of finite- and infinite-support distributions, and demonstrated that a simple decoder, performing the computation directly on the quantized measurements, achieves asymptotically equivalent performance to the fMMSE decoder. Although there are some technical restrictions on the source distributions and computations to ensure the high-resolution approximations are legitimate, the main goal of the paper is to show that DFSQ theory is widely applicable to distributed acquisition systems without requiring a complicated decoder. Furthermore, the asymptotic results give good approximations for the performance at moderate quantization rates.




Fig. 6. Empirical and theoretical performance are provided for the uniform, ordinary and functional quantizers for $N = 10$, exponential f_X , and $g(x_1^N) = \min(x_1^N)$. The distortions are multiplied by 2^{2R} to better indicate the convergence results.

DFSQ has immediate implications in how sensors in acquisition networks collect and compress data when the designer knows the computation to follow. Using both theory and examples, we demonstrate that knowledge of the computation may change the quantization mapping and improve fMSE. Because the setup is very general, there is potential for impact in areas of signal acquisition where quantization is traditionally considered as a black box. Examples include multi-modal imaging technologies such as 3D imaging and parallel MRI. This theory can also be useful in collecting information for applications in machine learning and data mining. In these fields, large amounts of data are collected but the measure of interest is usually some nonlinear, low-dimensional quantity. DFSQ provides insight on how data should be collected to provide more accurate results when the resources for acquiring and storing information are limited.

APPENDIX A PROOF OF THEOREM 1

Taylor's theorem states that a function g that is $n + 1$ times continuously differentiable on a closed interval $[a, x]$ takes the form

$$g(x) = g(a) + \left(\sum_{i=1}^n \frac{g^{(i)}(a)}{i!} (x-a)^i \right) + R_n(x, a),$$

with a Taylor remainder term

$$R_n(x, a) = \frac{g^{(n+1)}(\xi)}{(n+1)!} (x-a)^{n+1}$$

for some $\xi \in [a, x]$. More specific to our framework, for any $x \in [c_k, p_k]$, the first-order remainder is bounded as

$$|R_1(x, c_k)| \leq \frac{1}{2} \max_{\xi \in [c_k, p_k]} |g''(\xi)| (p_k - c_k)^2. \quad (31)$$

Using Condition UF2', we will uniformly bound $|g''(\xi)|$ by C_u .

The first mean-value theorem for integrals states that for a continuous function $r : [a, b] \rightarrow \mathbb{R}$ and integrable function $s : [a, b] \rightarrow [0, \infty)$ that does not change sign, there exists a value $x \in [a, b]$ such that

$$\int_a^b r(t)s(t) dt = r(x) \int_a^b s(t) dt. \quad (32)$$

For the case of the companding quantizers, combining this with (1) means

$$\frac{1}{K} = \int_{p_k}^{p_{k+1}} \lambda(x) dx = \lambda(y_k)(p_{k+1} - p_k) = \lambda(y_k)\Delta_k \quad (33)$$

for some $y_k \in (p_k, p_{k+1}]$, where we have defined the k th quantizer cell length $\Delta_k = p_{k+1} - p_k$. The relationship between K , λ , and Δ_k is central in the proof.

With these preparations, we continue to the proof. Consider expansion of $D_{\text{fmse}}(K, \lambda)$ by total expectation:

$$D_{\text{fmse}}(K, \lambda) = \sum_{k=0}^{K-1} \int_{p_k}^{p_{k+1}} |g(x) - g(c_k)|^2 f_X(x) dx. \quad (34)$$

We would like to eliminate the first and last terms of the sum because the unbounded interval of integration would cause problems with the approximation technique employed later. The last term is

$$\int_{p_{K-1}}^{\infty} |g(x) - g(p_{K-1})|^2 f_X(x) dx, \quad (35)$$

where we have used $c_K = p_{K-1}$. By Condition UF4', this is asymptotically negligible in comparison to

$$\left(\int_{p_{K-1}}^{\infty} \lambda(x) dx \right)^2 = \frac{1}{K^2}.$$

Thus the last term (35) does not contribute to $\lim_{K \rightarrow \infty} K^2 D_{\text{fmse}}(K, \lambda)$. We can similarly eliminate the first term, yielding

$$K^2 D_{\text{fmse}}(K, \lambda) \simeq K^2 \sum_{k=1}^{K-2} \int_{p_k}^{p_{k+1}} |g(x) - g(c_k)|^2 f_X(x) dx. \quad (36)$$

Effectively, Condition UF4' promises that the tail of the source distribution is decaying fast enough that we can ignore the distortion contributions outside the extremal codewords.

Further expansion of (36) using Taylor's theorem with remainder yields:

$$\begin{aligned} K^2 D_{\text{fmse}}(K, \lambda) &\simeq K^2 \sum_{k=1}^{K-2} \int_{p_k}^{p_{k+1}} |g'(c_k)(x - c_k) + R_1(x, c_k)|^2 f_X(x) dx \\ &= \underbrace{K^2 \sum_{k=1}^{K-2} \int_{p_k}^{p_{k+1}} |g'(c_k)|^2 |x - c_k|^2 f_X(x) dx}_{(A)} \\ &\quad + \underbrace{K^2 \sum_{k=1}^{K-2} 2 \int_{p_k}^{p_{k+1}} |R_1(x, c_k)| |g'(c_k)| |x - c_k| f_X(x) dx}_{(B)} \\ &\quad + \underbrace{K^2 \sum_{k=1}^{K-2} \int_{p_k}^{p_{k+1}} R_1(x, c_k)^2 f_X(x) dx}_{(C)}. \end{aligned}$$

Of the three terms, only term (A) has a meaningful contribution. It can be simplified as follows:

$$\begin{aligned} &K^2 \sum_{k=1}^{K-2} \int_{p_k}^{p_{k+1}} |g'(c_k)|^2 |x - c_k|^2 f_X(x) dx \\ &\stackrel{(a)}{=} K^2 \sum_{k=1}^{K-2} |g'(c_k)|^2 f_X(v_k) \int_{p_k}^{p_{k+1}} |x - c_k|^2 dx \\ &\stackrel{(b)}{=} K^2 \sum_{k=1}^{K-2} |g'(c_k)|^2 f_X(v_k) \frac{\Delta_k^3}{12} \\ &\stackrel{(c)}{=} \frac{1}{12} \sum_{k=1}^{K-2} f_X(v_k) \left(\frac{|g'(c_k)|^2}{\lambda^2(y_k)} \right) \Delta_k \\ &\stackrel{(d)}{\rightarrow} \frac{1}{12} \int_S \left(\frac{g'(x)}{\lambda(x)} \right)^2 f_X(x) dx, \end{aligned} \quad (37)$$

where (a) arises from using (32), where v_k is some point in the k th quantizer cell; (b) is evaluation of the integral, recalling (2); (c) follows from (33); and (d) holds as $K \rightarrow \infty$ by the convergence of Riemann rectangles to the integral (assumption UF3').

The higher-order error terms are negligible using the bound reviewed in (31). We now show that term (B) goes to zero:

$$\begin{aligned}
& K^2 \sum_{k=1}^{K-2} 2 \int_{p_k}^{p_{k+1}} |R_1(x, c_k)| |g'(c_k)| |x - c_k| f_X(x) dx \\
& \stackrel{(a)}{\leq} K^2 \sum_{k=1}^{K-2} C_u \Delta_k^2 |g'(c_k)| \int_{p_k}^{p_{k+1}} |x - c_k| f_X(x) dx \\
& \stackrel{(b)}{\leq} K^2 \sum_{k=1}^{K-2} C_u \Delta_k^4 |g'(c_k)| f_X(v_k) \\
& \stackrel{(c)}{=} \frac{C_u}{K} \sum_{k=1}^{K-2} f_X(v_k) \frac{|g'(c_k)|}{\lambda^3(y_k)} \Delta_k \\
& \stackrel{(d)}{\longrightarrow} 0,
\end{aligned} \tag{38}$$

where (a) follows from bounding $R_1(x, c_k)$ using (31); (b) arises from using (32) and bounding the integral; (c) follows from (33); and (d) holds as $K \rightarrow \infty$ by the convergence of Riemann rectangles to the integral (assumption UF3'). Hence, the distortion contribution becomes negligible as K increases.

A similar analysis can be used to show that expansion term (C) scales as $1/K^2$ with growing codebook size and is therefore also negligible.

APPENDIX B PROOF OF THEOREM 2

We parallel the proof of Theorem 1 using Taylor expansion and bounding the distortion contributions of each cell. We review the first-order version of the multivariate Taylor's theorem: a function that is twice continuously differentiable on a closed ball B takes the form

$$g(x_1^N) = g(a_1^N) + \sum_{n=1}^N [g_n(a_1^N)(x_n - a_n)] + R_1(x_1^N, a_1^N),$$

where we recall that $g_n(x_1^N)$ is the partial derivative of g with respect to the n th argument evaluated at the point x_1^N . The remainder term is bounded by

$$R_1(x_1^N, a_1^N) \leq \sum_{i=1}^N \sum_{j=1}^N (x_i - a_i)(x_j - a_j) C_m,$$

under Condition MF2'. Applying a linear approximation to a quantizer cell s with midpoint $(c_s)_1^N$ and side lengths $\{\Delta_i(s)\}_{i=1}^N$, the Taylor residual is

$$\begin{aligned}
|R_1(x_1^N, (c_s)_1^N)| & \leq C_m \sum_{i=1}^N \sum_{j=1}^N \Delta_i \Delta_j \\
& \stackrel{(a)}{\leq} N C_m \bar{\Delta} \sum_{i=1}^N \Delta_i \\
& \stackrel{(b)}{=} N C_m \bar{\Delta} \sum_{i=1}^N \frac{1}{\lambda_i(y_{s,i}) \kappa \alpha_i} \\
& \stackrel{(c)}{\leq} \frac{N C_m \bar{\Delta}}{\kappa \underline{\alpha}} \sum_{i=1}^N \lambda_i(y_{s,i})^{-1},
\end{aligned} \tag{39}$$

where in (a) we define $\bar{\Delta}$ as the longest quantizer interval length in any dimension; in (b) we invoke (33) with $y_{s,i}$ being the i th coordinate of some point in quantizer cell s ; and in (c) we define $\underline{\alpha}$ as the smallest α_i .

Let S_K be the partition lattice induced by N scalar quantizers, excluding the overload regions. By total expectation, we find the distortion of each partition cell and sum their contributions. By Condition MF4', the distortion from overload cells become negligible with increasing κ and can be ignored. Using Taylor's theorem, the scaled total distortion becomes

$$\kappa^2 D_{\text{fmse}}(K_1^N, \lambda_1^N) = A + B + C,$$

where

$$\begin{aligned}
A &= \kappa^2 \sum_{s \in \mathcal{S}_K} \int_{x_1^N \in s} \sum_{i=1}^N \sum_{j=1}^N g_i((c_s)_1^N) g_j((c_s)_1^N) \\
&\quad \cdot (x_i - c_{s,i})(x_j - c_{s,j}) f_{X_1^N}(x_1^N) dx_1^N, \\
B &= \kappa^2 \sum_{s \in \mathcal{S}_K} \int_{x_1^N \in s} 2 \sum_{n=1}^N |g_n((c_s)_1^N)| |x_n - c_{s,n}| \\
&\quad \cdot |R_1(x_1^N, (c_s)_1^N)| f_{X_1^N}(x_1^N) dx_1^N, \\
C &= \kappa^2 \sum_{s \in \mathcal{S}_K} \int_{x_1^N \in s} R_1^2(x_1^N, (c_s)_1^N) f_{X_1^N}(x_1^N) dx_1^N.
\end{aligned}$$

In term A , we may disregard all cross terms since $(X_n - c_{k,n})$ becomes uncorrelated in the high-resolution approximation because the pdf in each cell becomes well-approximated by a uniform distribution as the cell gets smaller. The remaining components of the distortion are

$$\kappa^2 \sum_{s \in \mathcal{S}_K} \int_{x_1^N \in s} \left(\sum_{n=1}^N g_n^2((c_s)_1^N) (x_n - c_{s,n})^2 \right) f_{X_1^N}(x_1^N) dx_1^N.$$

Using (33), the distortion contribution becomes

$$\kappa^2 \sum_{s \in \mathcal{S}_K} \int_{x_1^N \in s} \left(\sum_{n=1}^N \frac{g_n^2((c_s)_1^N)}{12K_n^2 \lambda_n^2(y_{s,n})} \right) f_{X_1^N}(x_1^N) dx_1^N,$$

where $y_{s,n}$ is the n th coordinate of some point in quantizer cell s . Using assumption MF3', this approaches the integral expression

$$\begin{aligned}
&\sum_{n=1}^N \frac{1}{12\alpha_n^2} \mathbb{E} \left[\left(\frac{g_n((c_s)_1^N)}{\lambda_n(X_n)} \right)^2 \right] \\
&= \sum_{n=1}^N \frac{1}{12\alpha_n^2} \mathbb{E} \left[\left(\frac{\gamma_n(X_n)}{\lambda_n(X_n)} \right)^2 \right].
\end{aligned}$$

It remains to show that the remainder terms B and C may be ignored. First, consider any of the N summands that constitute

B:

$$\begin{aligned}
& \kappa^2 \sum_{s \in \mathcal{S}_K} \int_{x_1^N \in s} 2|g_n((c_s)_1^N)| |x_n - c_{s,n}| \\
& \quad \cdot |R_1(x_1^N, (c_s)_1^N)| f_{X_1^N}(x_1^N) dx_1^N \\
& \stackrel{(a)}{\leq} \kappa^2 \sum_{s \in \mathcal{S}_K} \int_{x_1^N \in s} 2|g_n((c_s)_1^N)| |x_n - c_{s,n}| \\
& \quad \cdot \frac{NC_m \bar{\Delta}}{\underline{\alpha} \kappa} \sum_{i=1}^N \lambda_i(y_{s,i})^{-1} f_{X_1^N}(x_1^N) dx_1^N \\
& \stackrel{(b)}{\leq} \frac{2NC_m \bar{\Delta} \kappa}{\underline{\alpha}} \sum_{s \in \mathcal{S}_K} |g_n((c_s)_1^N)| \Delta_n(s) \\
& \quad \cdot \sum_{i=1}^N \lambda_i(y_{s,i})^{-1} f_{X_1^N}((v_s)_1^N) \prod_{j=1}^N \Delta_j(s) \\
& \stackrel{(c)}{\leq} \frac{2NC_m \bar{\Delta}}{\underline{\alpha}} \sum_{i=1}^N \sum_{s \in \mathcal{S}_K} |g_n((c_s)_1^N)| \\
& \quad \cdot \lambda_i(y_{s,i})^{-1} \lambda_n(y_{s,n})^{-1} f_{X_1^N}((v_s)_1^N) \prod_{j=1}^N \Delta_j(s) \\
& \stackrel{(d)}{\rightarrow} \left(\lim_{\kappa \rightarrow \infty} \frac{2NC_m \bar{\Delta}}{\underline{\alpha}} \right) \sum_{i=1}^N \mathbb{E} \left[\frac{|g_n(X_1^N)|}{\lambda_i(X_i) \lambda_n(X_n)} \right] \\
& = 0,
\end{aligned}$$

where (a) follows from (39); (b) employs (32) and bounds $|x_n - c_{s,n}| \leq \Delta_n$; (c) invokes (33); and (d) is valid according to assumption MF3'.

Remainder term C is negligible in a similar manner, which proves the theorem.

REFERENCES

- [1] D. L. Neuhoff, "The other asymptotic theory of lossy source coding," in *Coding and Quantization* (R. Calderbank, G. D. Forney, Jr., and N. Moayeri, eds.), vol. 14 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pp. 55–65, American Mathematical Society, 1993.
- [2] V. Misra, V. K. Goyal, and L. R. Varshney, "Distributed scalar quantization for computing: High-resolution analysis and extensions," *IEEE Trans. Inform. Theory*, vol. 57, pp. 5298–5325, Aug. 2011.
- [3] J. Z. Sun and V. K. Goyal, "Optimal quantization of random measurements in compressed sensing," in *Proc. IEEE Int. Symp. Inform. Theory*, (Seoul, Korea), pp. 6–10, June–July 2009.
- [4] J. Z. Sun and V. K. Goyal, "Scalar quantization for relative error," in *Proc. IEEE Data Compression Conf.*, (Snowbird, Utah), pp. 293–302, Mar. 2011.
- [5] M. Pugh and B. D. Rao, "Distributed quantization of order statistics with applications to CSI feedback," in *Proc. IEEE Data Compression Conf.*, (Snowbird, Utah), pp. 323–332, Mar. 2011.
- [6] D. Slepian and J. K. Wolf, "Noiseless coding of correlated information sources," *IEEE Trans. Inform. Theory*, vol. IT-19, pp. 471–480, July 1973.
- [7] A. B. Wagner, S. Tavildar, and P. Viswanath, "Rate region of the quadratic Gaussian two-terminal source-coding problem," *IEEE Trans. Inform. Theory*, vol. 54, pp. 1938–1961, May 2008.
- [8] R. Zamir and T. Berger, "Multiterminal source coding with high resolution," *IEEE Trans. Inform. Theory*, vol. 45, pp. 106–117, Jan. 1999.
- [9] A. D. Wyner and J. Ziv, "The rate-distortion function for source coding with side information at the decoder," *IEEE Trans. Inform. Theory*, vol. IT-22, pp. 1–10, Jan. 1976.
- [10] A. Orlitsky and J. R. Roche, "Coding for computing," *IEEE Trans. Inform. Theory*, vol. 47, pp. 903–917, Mar. 2001.
- [11] T. S. Han and K. Kobayashi, "A dichotomy of functions $F(X, Y)$ of correlated sources (X, Y) from the viewpoint of the achievable rate region," *IEEE Trans. Inform. Theory*, vol. IT-33, pp. 69–76, Jan. 1987.
- [12] V. Doshi, D. Shah, M. Médard, and S. Jaggi, "Distributed functional compression through graph coloring," in *Proc. IEEE Data Compression Conf.*, (Snowbird, Utah), pp. 93–102, Mar. 2007.
- [13] H. Yamamoto, "Wyner–Ziv theory for a general function of the correlated sources," *IEEE Trans. Inform. Theory*, vol. IT-28, pp. 803–807, Sept. 1982.
- [14] H. Feng, M. Effros, and S. A. Savari, "Functional source coding for networks with receiver side information," in *Proc. 42nd Annu. Allerton Conf. Commun. Control Comput.*, pp. 1419–1427, Sept. 2004.
- [15] H. V. Poor, "High-rate vector quantization for detection," *IEEE Trans. Inform. Theory*, vol. 34, pp. 960–972, Sept. 1988.
- [16] G. R. Benitz and J. A. Bucklew, "Asymptotically optimal quantizers for detection of i.i.d. data," *IEEE Trans. Inform. Theory*, vol. 35, pp. 316–325, Mar. 1989.
- [17] R. Gupta and A. O. Hero, III, "High-rate vector quantization for detection," *IEEE Trans. Inform. Theory*, vol. 49, pp. 1951–1969, Aug. 2003.
- [18] J. A. Bucklew, "Multidimensional digitization of data followed by a mapping," *IEEE Trans. Inform. Theory*, vol. IT-30, pp. 107–110, Jan. 1984.
- [19] T. Linder, R. Zamir, and K. Zeger, "High-resolution source coding for non-difference distortion measures: Multidimensional companding," *IEEE Trans. Inform. Theory*, vol. 45, pp. 548–561, Mar. 1999.
- [20] W. R. Bennett, "Spectra of quantized signals," *Bell Syst. Tech. J.*, vol. 27, pp. 446–472, July 1948.
- [21] P. F. Panter and W. Dite, "Quantizing distortion in pulse-count modulation with nonuniform spacing of levels," *Proc. IRE*, vol. 39, pp. 44–48, Jan. 1951.
- [22] J. A. Bucklew and G. L. Wise, "Multidimensional asymptotic quantization theory with r th power distortion measures," *IEEE Trans. Inform. Theory*, vol. IT-28, pp. 239–247, Mar. 1982.

- [23] S. Cambanis and N. L. Gerr, "A simple class of asymptotically optimal quantizers," *IEEE Trans. Inform. Theory*, vol. IT-29, pp. 664–676, Sept. 1983.
- [24] T. Linder, "On asymptotically optimal companding quantization," *Prob. Contr. Inform. Theory*, vol. 20, no. 6, pp. 475–484, 1991.
- [25] V. K. Goyal, "High-rate transform coding: How high is high, and does it matter?," in *Proc. IEEE Int. Symp. Inform. Theory*, (Sorrento, Italy), p. 207, June 2000.
- [26] R. M. Gray and A. H. Gray, Jr., "Asymptotically optimal quantizers," *IEEE Trans. Inform. Theory*, vol. IT-23, pp. 143–144, Feb. 1977.
- [27] A. György, T. Linder, P. A. Chou, and B. J. Betts, "Do optimal entropy-constrained quantizers have a finite or infinite number of codewords?," *IEEE Trans. Inform. Theory*, vol. 49, pp. 3031–3037, Nov. 2003.
- [28] H. Gish and J. P. Pierce, "Asymptotically efficient quantizing," *IEEE Trans. Inform. Theory*, vol. IT-14, pp. 676–683, Sept. 1968.
- [29] R. M. Gray and D. L. Neuhoff, "Quantization," *IEEE Trans. Inform. Theory*, vol. 44, pp. 2325–2383, Oct. 1998.
- [30] A. Segall, "Bit allocation and encoding for vector sources," *IEEE Trans. Inform. Theory*, vol. IT-22, pp. 162–169, Mar. 1976.
- [31] D. Hui and D. L. Neuhoff, "Asymptotic analysis of optimal fixed-rate uniform scalar quantization," *IEEE Trans. Inform. Theory*, vol. 47, pp. 957–977, Mar. 2001.